arXiv:1603.01324v1 [cs.RO] 4 Mar 2016

Compressed Sensing for Tactile Skins

Brayden Hollis, Stacy Patterson, and Jeff Trinkle

Abstract— **Whole body tactile perception via tactile skins offers large benefits for robots in unstructured environments. To fully realize this benefit, tactile systems must support real-time data acquisition over a massive number of tactile sensor elements. We present a novel approach for scalable tactile data acquisition using compressed sensing. We first demonstrate that the tactile data is amenable to compressed sensing techniques. We then develop a solution for fast data sampling, compression, and reconstruction that is suited for tactile system hardware and has potential for reducing the wiring complexity. Finally, we evaluate the performance of our technique on simulated tactile sensor networks. Our evaluations show that compressed sensing, with a compression ratio of 3 to 1, can achieve higher signal acquisition accuracy than full data acquisition of noisy sensor data.**

I. INTRODUCTION

With the advances in robot technology, there is an increased desire to move robots from controlled environments to unstructured human environments. The National Robotics Initiative states that part of its goal is "to accelerate the development and use of robots in the United States that work beside or cooperatively with people" [1]. To achieve this goal, robots require improved perception to safely handle the uncertain and dynamic structure of these environments. Improved perception becomes even more important as robots need to effectively manipulate items in the environment. Tactile sensors, sensors that gather data through contact (*e.g.*, force and temperature), are an active research area with high potential to help address this issue.

An important area of research in tactile sensing is tactile systems that cover large portions of robot bodies. These systems, often called *tactile skins*, potentially provide improvements in a robot's awareness and manipulation abilities. For example, Ohmura et al. [2] demonstrate a robot lifting a heavy box with use of a tactile skin, a task beyond the robot's capability with traditional methods. A number of other tasks, especially for complex robots like humanoids, would benefit from tactile sensors on the surfaces of various links.

One of the challenges with developing tactile skins is the data acquisition process. It has been suggested that to achieve human-like perception, tactile systems should have spatial resolution up to 1 mm [3]. This translates to on the order of tens of thousands to millions of individual sensing elements, known as taxels, for human adult-sized robots (currently the largest system contains approximately 4200 taxels [4]). Full utilization of the data involves integrating it into control loops. This requires real-time acquisition and processing rates on the order of 1 kHz (*e.g.*, for high-fidelity force control). In addition, due to the distributed nature of the system, a physical network is required for data transfer but is limited by physical requirements (*e.g.*, weight, space, etc.) [3]. Finally, sensors are typically noisy, which can further affect the acquisition process. These combined requirements and constraints pose a challenge to acquiring the data from tactile skins.

There have been multiple proposed and implemented tactile skins from the research community. Most use a data acquisition process, which we refer to as full data acquisition, of polling all the sensors directly [5], [6], [7], relying on fast hardware for high sampling rates. These systems typically have a subset of local sensors that are read serially via a micro-controller which then passes the data on a high capacity network it shares with other micro-controllers. The number of sensors that can be associated with a micro-controller is limited. The number of micro-controllers connected to the network is also limited to maintain the desirable high sampling rate. To scale this approach to full body tactile skins will require multiple of these high capacity networks, which is considered undesirable [3] due to physical constraints. This method has been implemented with a sampling rate of up to 100 hz (though the number of sensors is unclear as this value is given for the local level sampling) [6].

An alternate approach to full data acquisition is to use local processing of sensor data within small clusters of sensors [8], [6]. This approach speeds up the entire tactile data processing rates through parallel local processing and only transferring aggregate information. Local processing like this suffers from a loss in data quality or a reduction in the versatility of the system as features must be determined *a priori*. Full data acquisition, on the other hand, has the freedom to extract the various relevant features for many different applications. Local processing has been implemented with sampling rates up to 500 hz [6].

We propose to address the tactile data acquisition problem using compressed sensing [9]. Compressed sensing is a popular technique for signal processing in which the signal is simultaneously sampled and compressed. Provided the original signal is sparse in some basis, it can be efficiently

This work was partially supported by NSF Grant CCF-12084 (through the National Robotics Initiative), NSF Grant CNS-1527287, and the NSF Independent Research and Development Program. B. Hollis is supported by a SMART Scholarship, funded by OASD/R&E (Office of the Assistant Secretary Defense-Research and Evaluation), Defense Wide / PE0601120D8Z National Defense Education Program (NDEP) / BA-1, Basic Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

B. Hollis, S. Patterson, and J. Trinkle are with the Department of Computer Science, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY, USA, hollib@rpi.edu,{sep,trink}@cs.rpi.edu.

reconstructed from this compressed version with little to no loss. In this work, we demonstrate the feasibility of a compressed sensing-based tactile data acquisition approach using existing compressed sensing tools. We first show that tactile sensor data is amenable to compressed sensing by identifying a basis under which the data is sparse. We then identify sampling and compression operators that are suited for the tactile system hardware. Finally, we explore different methods for real-time signal reconstruction from the compressed data.

Our approach shares the benefits of both the approaches previously discussed and overcomes the short comings. Similar to the local processing approach, our approach reduces the amount of data to be transferred, allowing for faster sampling rates, but it also allows the system to have the full data set and the resulting flexibility for multiple applications, like the full data acquisition approach. Our work also has potential to reduce wiring complexity as the sampling and compression is amiable to hardware implementations that share wires across multiple taxels.

The rest of the paper is organized as follows. Section II describes the particular problem we address and the data set we used. We then describe compressed sensing and compress sensing requirements related to tactile skins in Section III. Our proposed solution is discussed in Section IV and is evaluated in Section V. The paper then finishes with concluding remarks and future work in Section VI.

II. PROBLEM SETTING

A. Problem Description

We address the problem of real-time data acquisition from a large-scale array of taxels. The array consists of N individual taxels arranged in a $\sqrt{N} \times \sqrt{N}$ flat grid. In this paper, we consider tactile skins whose elements each output the magnitude of the contact force normal to the plane of the array at its location. However, we believe our approach can be extended to other tactile sensors that measure one-dimensional data, such as temperature.

We represent the sensor data as a vector $x(t) \in \mathbb{R}^N$, where each element, $x_i(t)$, corresponds to the force applied to i^{th} taxel at time-step t. The sensors do not measure the applied force exactly, but rather, they measure $x(t) + \epsilon(t)$, where $\epsilon(t) \in \mathbb{R}^N$ is the sensor noise. The tactile data acquisition problem is to efficiently transfer $x(t)$ (or a close approximation of it) from the sensors to the central processing unit. We note that "efficiently" is a subjective term that depends on the system and application. Ideally, we aim for solutions that require on the order of milliseconds so that the sensed data can be used in real-time feedback control. Additionally, we desire solutions that reduce the network hardware complexities.

As discussed in the introduction, approaches that directly collect sensed data suffer from scalability limitations. We consider an alternate approach based on the new and popular technique of compressed sensing [9]. In compressed sensing, the data is compressed in hardware during the acquisition process. This compressed data can be more efficiently transmitted to the central processing unit, since the reduced data requires fewer high capacity networks. At the central processing unit, the original signal can be efficiently recovered. This unique approach will enable scaling tactile skins to large numbers of sensors due to the data reduction and the hardware compression techniques that allow simpler wiring.

B. Data Set Description

For our study, we used the Gazebo 3.0 [10] simulator with the SkinSim 0.1.0 simulation environment [11] to generate data sets. SkinSim 0.1.0 models a tactile skin as a layer of spheres individually attached to a surface by a spring and damper. SkinSim is currently limited to planar grid arrays. We set up our system with each sphere-spring-damper system representing a single taxel with a sphere radius of 2.5 mm. We had no space between taxels. We generated sensor data using two different sensor array sizes, 40×40 and 64×64, which matches the number of taxels in the current largest system. Each sensor measured contact normal force values between 0 and 2.5 Newtons.

We use four different scenarios to demonstrate that our approach works with contact at different scales (*i.e.*, number of active sensors), shape complexities, and locations on the grid. In the first scenario, LIFT BOX, we used the 10 cm wooden cube from Gazebo's default models. We dropped the object onto the sensor pad, with the face parallel to the pad. After the object settled, it was lifted from the pad with constant velocity of 0.1 $\frac{\text{m}}{\text{s}}$. The second scenario, LIFT WRENCH, followed the same procedure, but used Gazebo's default model for a monkey wrench. In the third scenario, PATH, a metal peg, which was forced to remain vertical throughout the scenario, was dropped in a random position on the sensor pad and then moved along a random path parallel to the plane of the tactile array. The final scenario, DRAG WRENCH, again used the monkey wrench. After dropping onto the sensor pad, the wrench was moved with constant velocity of 0.05 $\frac{\text{m}}{\text{s}}$ also parallel to the sensor plane. The sensors' heights caused them to rest above the ground plane, so, as the wrench moved off the sensor pad, it was partially supported by the lower ground plane. All of the scenarios except PATH were run until the object stopped making contact with the sensor array. PATH had no guarantees of leaving the sensor pad, so it was run for an extended period of time.

For the experiments on the 40×40 and 64×64 arrays, the simulation time-steps were chosen to be 1 and 0.1 milliseconds respectively. Table I shows the number of time-steps in which we collected data for each scenario. We set high sampling rates as dynamics simulators perform better with small time-steps. We used all the sampling data since we are more interested in a sequence of tactile data than the specific scenario. So, while the simulator was set up with the object falling under gravity, due to realistic tactile sampling rates, we treat the scenario more as the object being manipulated against the sensors with a slower velocity. Even with the small time-steps, the data generated by the SkinSim

simulator exhibited high-frequency oscillations in the contact force signals. We believe this is due to integration error of the physics engine or to low damping. We compensated for this physically unrealistic behavior by applying a moving average filter with a width of 10 time-steps.

TABLE I: Number of simulation time-steps of data for each scenario

Scenario	Array	number of time-steps
LIFT BOX	40×40	700
LIFT WRENCH	40×40	3000
DRAG WRENCH	40×40	5000
PATH	40×40	12200
LIFT BOX	64×64	7600
DRAG WRENCH	64×64	12600

III. BACKGROUND ON COMPRESSED SENSING

A. Compressed Sensing Theory

Let $x \in \mathbb{R}^N$ be the signal of interest, for example, the force readings of the tactile sensors at a given time. In compressed sensing, the signal is compressed by taking M linear measurements, each of the form,

$$y_i = \sum_{j=1}^{N} \Phi_{ij} x_j.$$

The *measurement vector* $y \in \mathbb{R}^M$, with components y_i, is given by

$$y = \Phi x, \tag{1}$$

where $\Phi = [\Phi_{ij}]$ is the $M \times N$ *measurement matrix*.

If $M < N$, the linear system (1) is under-determined, and thus, in general, it is not possible to recover x from fewer than N measurements. The theory of compressed sensing shows that, if x is sparse, meaning it has only a few non-zero components, and Φ satisfies a property known as the *restricted isometry property* [9], the vector x can be recovered from $M << N$ measurements, on the order of $O(\log N)$ under ideal conditions. Typically, the signal x is not itself sparse, but may be sparse in some *representation basis*. Specifically, a signal x is sparse in a basis $\Psi \in \mathbb{R}^{N \times N}$ if there is a sparse vector θ such that $x = \Psi\theta$. In this case, the measurement vector can be written as

$$y = \Phi\Psi\theta.$$

The objective is then to recover θ from the measurements. The compressed sensing recovery problem can be formalized as

$$\underset{\theta \in \mathbb{R}^N}{\text{minimize}} \ \|\theta\|_0 \quad \text{subject to } \Phi\Psi\theta = y \tag{2}$$

Here $\|\cdot\|_0$ denotes the ℓ_0 pseudonorm, *i.e.*, the number of non-zero components. Once the solution to this problem $\hat{\theta}$ has been obtained, one can recover the original non-sparse signal x as $x = \Psi\hat{\theta}$.

In this work, we consider a variation on the compressed sensing problem where measurements may not be exact. The measurement vector is given by

$$y = \Phi\Psi\theta + \epsilon,$$

where $\epsilon \in \mathbb{R}^M$ is the measurement noise. This noise can model both errors in the individual sensor readings as well as in the measurement acquisition. The sparse vector θ can be found as the solution of the following optimization problem

$$\underset{\theta \in \mathbb{R}^N}{\text{minimize}} \ \tfrac{1}{2}\|\Phi\Psi\theta - y\ \|_2^2 + \lambda\|\theta\|_0, \tag{3}$$

where λ is a scalar-valued parameter used to tune the level of sparsity of the solution. It is intractable to solve problem (3) directly, however, efficient alternate approaches to find $\hat{\theta}$ have been derived [9].

The approach we use in this work is Basis Pursuit Denoising (BPDN). In BPDN, one solves a convex relaxation of (3),

$$\underset{\theta \in \mathbb{R}^N}{\text{minimize}} \ \tfrac{1}{2}\|\Phi\Psi\theta - y\ \|_2^2 + \lambda\|\theta\|_1. \tag{4}$$

Since this problem is convex, it can be solved efficiently using one of many algorithms for convex optimization, as well as variants developed specifically for compressed sensing. It has been shown that BPDN produces optimal or near-optimal sparse solutions in a wide variety of settings [12].

B. Compressed Sensing for Tactile Skins

To apply compressed sensing to tactile skins, we want to take linear combinations of subsets of taxels' values to generate measurements. From those measurements, we want to reconstruct the individual taxels' values. This is done once per time-step and allows us to recover the full data while reducing the number of measurements.

To successfully apply compressed sensing to the tactile sensing problem, we must address the following three components:

- **Identification of a representation basis.** We must identify a basis under which the signals (sensor data) from our data sets are sufficiently sparse. A greater level of sparsity means that the signals can be accurately reconstructed from fewer measurements. Further, the representation basis should be universal, meaning that it can be used to "sparsify" signals in all scenarios.
- **Identification of a sparse measurement matrix that is compatible with the basis and hardware.** As mentioned above, it is necessary that $\Phi\Psi$ satisfy the restricted isometry property. In the tactile sensing application, we have a further restriction that the measurement matrix be sparse. This sparsity is critical as it allows the measurement matrix to be implemented in hardware and reduce the wiring complexity. This is shown in Figure 1, where in Figure 1a shows each signal with a dedicated connection and Figure 1b shows wiring designed from a compressed sensing measurement matrix where sets of taxels are daisy-chained. Sparsity is necessary since the non-zero elements in each row indicate the linear

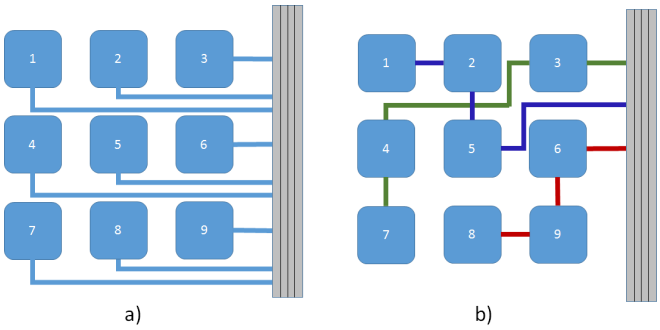

Fig. 1: Example wiring schematics for a) individual sensor measurements and b) compressed sensing aggregated measurements on a 3×3 tactile grid.

combinations (*i.e.*, the taxels with corresponding non-zero values in that row need to be part of the same daisy-chain, such as taxels 1, 2, and 5 in Figure 1b) and the more taxels per linear combination typically corresponds to longer and more complex daisy-chains. So we seek to minimize this interaction to reduce the amount of wiring.

- **Selection/implementation of a fast signal reconstruction algorithm.** The reconstruction algorithm needs to be able to reconstruct the signal in real time so the signal can be used in control loops. As reconstruction is a tradeoff between time and accuracy, it also ought to reconstruct the signal at least as a good approximation of the noisy signal, but ideally as a good approximation of the actual signal.

In the following section, we describe our approach to each of these three components.

IV. SOLUTION

A. Representation Basis

Historically, tactile processing often used computer vision techniques since the common array pattern of tactile sensor systems corresponds well with image pixel arrays [13]. Following this trend, we looked at compressed sensing of images to identify a basis Ψ that will render the tactile data signal sufficiently sparse. As candidate bases, we considered the Daubechies wavelet transforms, which are similar to the wavelet transforms in JPEG2000. Daubechies wavelet transforms are a set of transforms where each individual transform is referenced by an even number (*e.g.*, D4) such that larger numbers correspond to more complex transformations [14].

We choose to use the D2 wavelet basis as, for all scenarios, the D2 wavelet (also known as the Haar wavelet) basis outperformed the D4 wavelet basis, yielding signals that were indeed sparse. This can be seen in Table II, which shows the number of non-zero components in the vector $\Psi x(t)$, both when Ψ is the D2 wavelet and when Ψ is the D4 wavelet. The table shows the mean and maximum number of non-zero components over all time-steps for each scenario.

One point to note is that, while we represent the signal as a one-dimensional vector, the sensor array is a two-dimensional structure. Thus we use the two-dimensional Haar wavelet for our basis in the same manner that it is used for two-dimensional images.

TABLE II: The number of non-zero elements over time for different bases.

Scenario	N	D2		D4	
		mean	max	mean	max
LIFT BOX	1600	108.3	266	750.1	1575
LIFT WRENCH	1600	92.2	388	282.2	1575
DRAG WRENCH	1600	84.9	449	254.3	1575
PATH	1600	30.3	266	139.8	1575
LIFT BOX	4096	189.8	263	929.1	961
DRAG WRENCH	4096	240.3	670	602.7	1208

B. Sensing Matrix

As stated in Section III-A, it is necessary that the sensing matrix Φ be chosen so that $\Phi\Psi$ satisfies the restricted isometry property. Unfortunately, it is NP-Hard to verify whether a given matrix $\Phi\Psi$ satisfies this property. It has been shown that if Φ is a random matrix with entries drawn from a Gaussian distribution, the restricted isometry property is satisfied with high probability [15]. While this choice of Φ is feasible from an algorithmic perspective, such a matrix is dense. Thus, it is problematic to implement this measurement approach in hardware due to the issues discussed in Section III-B.

We instead select a sparse measurement matrix that will support data collection and compression with less wiring. We use the Scrambled Block Hadamard Ensemble (SBHE) developed by Gan et al. [16]. It has been shown that for this choice of Φ, the product $\Phi\Psi$ behaves like a random Gaussian matrix for a wide variety of bases. SBHE is a partial block Hadamard transform with randomly permuted columns. It can be represented as

$$\Phi_H = Q_M W P_N, \tag{5}$$

where W is a $N \times N$ block diagonal matrix with each block a $B \times B$ Hadamard matrix, P_N is the permutation matrix, which randomly reorders the N columns of W, and Q_M uniformly at random selects M rows of WP_N. As the non-zero elements come from the $B \times B$ Hadamard blocks in W, B determines the number of sensors connected together for each measurement element. From a wiring standpoint, we want to minimize B and, similar to Gan et al., we found $B = 32$ worked well for our evaluations.

SBHE also has two additional benefits. First, all non-zero elements are either $+1$ or -1. This simplifies the hardware as taxels only need the additional hardware to negate their values instead of various hardware for different gains for each measurement, as would be the case for the random Gaussian matrix. Second, the taxels are divided into disjoint measurement subsets that each give multiple measurements. This means for any two measurements, either the whole subset of taxels or none of the taxels are used in both measurements. This is a benefit as no wires are needed between the subsets, allowing for very distinct daisy-chains.

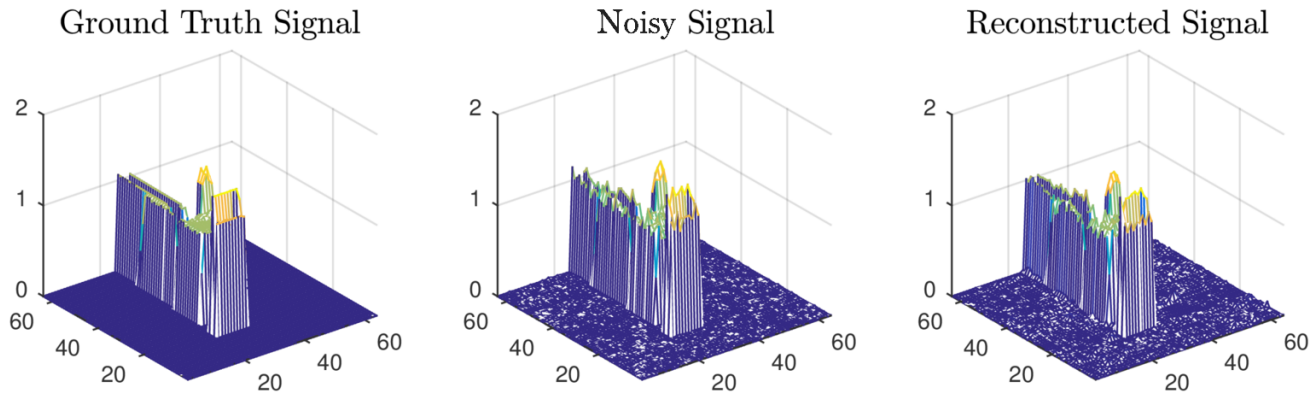

Fig. 2: Sensor values from DRAG WRENCH (64×64) scenario for the ground truth signal (left), noisy signal (middle) and the reconstructed signal (right). The reconstruction was from 1365 measurements using 20 iterations of the FISTA algorithm and has a PSNR of 31.2.

TABLE III: Quality of Signal Reconstruction with Various Number of Measurements using 30 Iterations

Scenario (No. of Elements)	No. of Measurements	Reconstructed Signal Mean PSNR	Reconstructed Signal PSNR Range	Noisy Signal Mean PSNR	Noisy Signal PSNR Range
LIFT BOX (1600)	400	52.8	39.5 - 63.2	46.5	39.1 - 48.7
	533	52.5	40.9 - 60.8		
	800	52.1	41.4 - 57.6		
LIFT WRENCH (1600)	400	52.1	26.5 - 62.3	47.2	38.4 - 48.8
	533	52.7	31.1 - 60.4		
	800	52.6	36.9 - 58.0		
DRAG WRENCH (1600)	400	44.7	25.3 - 61.9	46.4	37.4 - 48.7
	533	47.7	28.5 - 59.8		
	800	48.9	34.5 - 58.0		
PATH (1600)	400	56.3	45.4 - 63.1	48.0	44.8 - 49.1
	533	55.9	49.2 - 60.7		
	800	55.0	50.3 - 58.4		
LIFT BOX (4096)	1024	49.8	38.6 - 54.9	46.8	36.6 - 48.5
	1365	49.9	39.0 - 54.1		
	2048	50.2	38.3 - 53.6		
DRAG WRENCH (4096)	1024	45.4	23.3 - 54.8	46.4	36.2 - 48.3
	1365	46.4	29.8 - 53.9		
	2048	48.3	37.0 - 53.2		

This property also has the potential for the measurements to share wires.

C. Reconstruction Algorithm

To solve problem (4), we use an iterative, gradient-like method called Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [17]. We selected this algorithm because of its fast convergence; FISTA converges at a rate of $O(1/k^2)$, where k is the iteration number, whereas standard gradient methods converge at a rate of $O(1/k)$. Every time-step we use FISTA to reconstruct the signal from the measurements taken in that time-step. As the signal typically does not change much between time-steps, we use the previous time-step's reconstructed signal as the initial estimate for the current signal and use FISTA to refine the estimate into a solution. This warm-start led to faster convergence to the solution. FISTA is also amenable to hardware-based implementations, which can further accelerate its performance. We exploit this feature in our implementation.

V. SIMULATION RESULTS

A. Algorithm Implementation

We implemented FISTA in Matlab 2015a, utilizing Matlab's GPU capabilities to speed up the algorithm's execution. The resulting code required approximately 1ms per iteration (see Table IV). We ran the algorithm for a constant number of iterations for each compressed sensing problem as it standardized the reconstruction time, and checking for other possible exit criteria greatly reduced performance. We performed simulations for several values of λ (see (4)) and determined that $\lambda = 0.1$ worked well for all scenarios. We used this parameter value in all simulations.

We noted that the reconstructed signal returned from FISTA sometimes contained negative values, but as we know

this is not a feasible signal, we set those elements to 0.

Our experiments were run on a machine using Linux Ubuntu 14.04. The machine contained an Intel Core i7 CPU (8 cores at 3.07 GHz) and a single GeForce GTX 480 GPU.

B. Noise Model

We added noise to the tactile data before performing the measurements. Our noise model was a normal distribution with zero mean. The model had a standard deviation of 0.0628 Newtons (5 percent) at the sensor range midpoint and tapered towards the ends of the sensor range. All values beyond the sensor range were modified to the corresponding boundary value.

C. Results

To test our method, we applied our measurement matrix Φ_H to our scenarios and then reconstructed the signal with FISTA. An illustrative example of the sensor data and reconstruction is shown in Figure 2. The figure shows the three signals for a single time-step: the ground truth signal (*i.e.*, the SkinSim filtered signal) of the object contact, the noisy signal (*i.e.*, the ground truth signal plus our noise model) that we compress for our measurements, and the reconstructed signal (*i.e.*, the restored signal from FISTA). From the figure it is evident that the reconstructed signal closely matches the shape and magnitude of the original signal, demonstrating our method gives reasonable results.

To evaluate the quality of our reconstruction, for each time-step we measured the peak signal-to-noise ratio (PSNR) of the reconstructed signal and compared it to the PSNR of the noisy signal (larger PSNRs are better).

We first evaluated the impact of the number of measurements on the reconstruction quality. We varied the number of measurements to $\frac{N}{4}$, $\frac{N}{3}$ and $\frac{N}{2}$, where N is the number of taxels. For these experiments, the number of FISTA iterations was set to 30. We selected 30 iterations after tracking the number of iterations FISTA ran when using various stopping criteria based on convergence (*e.g.*, the change in the solution between iterations was below a threshold, solution was within a threshold of the ground truth, etc.) and found 30 iterations was typically more than enough. The mean and range of the PSNR over time are shown in Table III for each scenario, along with the PSNR of the noisy signal. The table shows that for the vast majority of the reconstructions, the mean PSNR exceeded that of the noisy signal. This means that the reconstructed signal more accurately estimated the ground truth than the noisy signal. In addition, for most scenarios, the reconstructed mean PSNR was similar for all three measurement sizes, though the PSNR range was typically larger for fewer measurements. This means that while more measurements made the reconstruction quality more consistent, on average the PSNR were independent of the number of measurements. The scenarios that have mean PSNR less than the noisy signal both used only $\frac{N}{4}$ measurements. This indicates that $\frac{N}{4}$ measurements can be comparable with the noisy signal accuracy only for some scenarios.

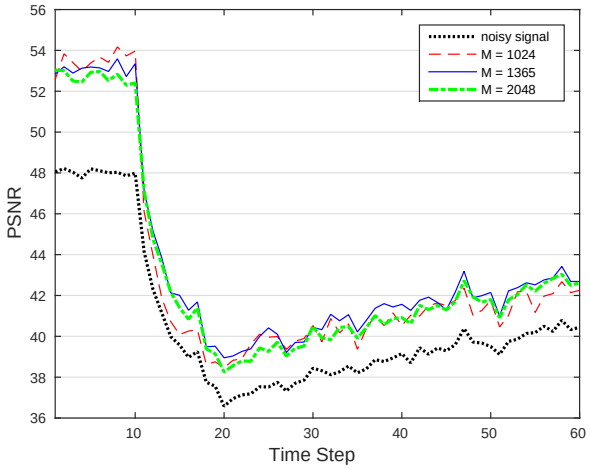

(a) Representative Case Time Sequence - LIFT BOX 64×64

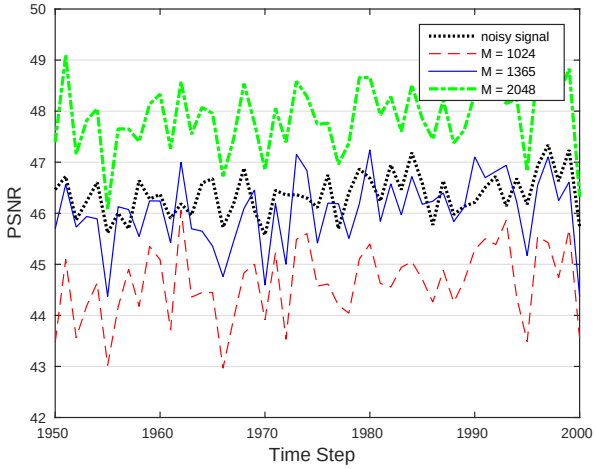

(b) Worst Case Time Sequence - DRAG WRENCH 64×64

Fig. 3: Representative time sequences of PSNR of the reconstructed signal for the different numbers of measurements ($\frac{N}{4}$, $\frac{N}{3}$, $\frac{N}{2}$) as well as the PSNR of the noisy signal. (a) is representative of most scenarios as the reconstructed signal outperforms the noisy signal while (b) is the worst-case scenario where the reconstruction signal only out performs the noisy signal for large numbers of measurements.

The lower end of the PSNR range values occurred during the first few time-steps of contact between the object and sensor pad. This visually results in a dip in the PSNR plot as seen in Figure 3a. The reason for this dip is our noise model produces more noise at the middle of the sensor range than at the ends, and prior to contact the sensor values are at the low end of the range and during first contact, the sensor values are closer to the middle of the sensor range.

Figure 3 shows representative time sequences of how the number of measurements affects the reconstruction using the LIFT BOX and DRAG WRENCH scenarios, both using the 64×64 array. The results for the LIFT BOX scenario in

Figure 3a are fairly representative of most of the scenarios. The reconstructed PSNR does not vary much for different numbers of measurements and is larger than the noisy PSNR. Figure 3b shows a representative snapshot of the worst performing scenarios, the DRAG WRENCH scenarios. The PSNR for $\frac{N}{4}$ measurements was consistently below that of the noisy signal, the PSNR for $\frac{N}{3}$ measurements was similar to that of the noisy signal, and for $\frac{N}{2}$ measurements, the PSNR was better than that of the noisy signal. This figure shows that, in the worst case, using $\frac{N}{3}$ measurements gives accuracy on par with the noisy signal. These results show that our measurement approach compresses the signal up to a third of its original size, with the ability to recover the original signal typically with higher quality than the noisy signal.

We next looked at the impact of the number of iterations on the reconstruction accuracy, which is an indication of the reconstruction speed. We ran FISTA for 10, 20, and 30 iterations and evaluated the quality of the reconstructed signals. We used $\frac{N}{3}$ measurements in all cases. The results are presented in Table IV. The average PSNR for each scenario was unaffected by the number of iterations, but the minimum PSNR varied for most scenarios, especially for 10 iterations. This is especially of interest as the few time-steps that produced the lower end of the PSNRs' range were the only time steps it was challenging for the reconstructed signal to be comparable to the noisy signal.

To further illustrate how the number of iterations affects the PSNR, Figure 4 shows representative time sequences for the parts of the reconstructed signal using scenario LIFT BOX 64×64. Both Figures 4a and 4b show the PSNR for reconstructed signal after 10, 20, and 30 FISTA iterations, along with the PSNR for the noisy signal for comparison. Figure 4a is a snapshot of the majority of the time sequences where the ground truth signal changes slowly, while Figure 4b covers the time sequence of the first contact. Where the ground truth signal changes slowly, which is the majority of the time for all scenarios, the reconstructed signals' PSNR plots are nearly indistinguishable for the different iterations, which is shown in Figure 4a. Figure 4b similarly shows the PSNR plots are nearly identical, except that for a few time-steps the PSNR for 10 iterations was significantly lower than for 20 and 30 iterations. These time-steps occurred when the object initially made contact with the sensors, or in other words, the deviation between the PSNR of 10 iterations and the other values occurred at sudden, significant signal changes. The drastic signal change requires more iterations to converge as the warm-start is no longer a good initial estimate, and Figure 4b, along with Table IV, shows 10 iterations is not enough. The PSNR stabilizes again due to the aggregation of FISTA iterations, which occurs due to the warm-starts. While for some applications, the loss in quality that 10 iterations experienced might be acceptable, this is not always the case. Significant signal changes often indicate significant changes in the robot state and thus, an accurate reconstruction may be important for generating the correct response. In this case, using 20 iterations is preferable.

VI. CONCLUSION AND FUTURE WORK

A. Conclusion

We have proposed an approach to tactile data acquisition using compressed sensing. In our approach, tactile data is compressed in hardware before transmission to the central processing unit. Then, the fully sensed data signal is reconstructed from the compressed data.

In this work, we have used standard (*i.e.*, pre-existing and non-optimized) tools of compressed sensing, and we still achieved full signal reconstruction of 4096 taxels at a rate of 50 hz, which is on the same order of magnitude as the sensors presented by Schmitz et al. [6], the sensors used for the system with the most taxels. Further, our system was able to compress the signal to a third of its original size and produce a higher quality signal than the noisy signal a vast majority of the time. Additionally, we have discussed how compressed sensing may provide guidance for new wiring techniques with potential to reduce the wiring complexity. Using daisy-chain techniques, wiring could be reduced on the same order as the data compression. We were able to reduce the number of measurements by a factor of three, but the theoretical limit for compressed sensing is to have $O(\log N)$ measurements. From this initial investigation, we conclude that compressed sensing is a feasible approach for tactile data acquisition and worth further exploration. We believe this paradigm can open the door to larger-scale tactile systems as well as faster data acquisition.

B. Future Work

We plan to continue to develop our system so as to meet the suggested requirements mentioned in the introduction by exploring additional bases and measurement matrices. Not only will we look for representation bases that give even more sparsity in the tactile data, which is a key component for achieving $O(\log N)$ measurements, but we also will explore bases that work with different tactile sensor arrangements (*e.g.*, non-planar). Part of our future work in exploring measurement matrices will be to consider the constraint that sensors that interact during measurement be in close proximity. This will further simplify the wiring arrangement. Finally, we will investigate alternatives and enhancements to our reconstruction algorithm to improve reconstruction speed.

REFERENCES

[1] "National robotics initiative, program solicitation," http://www.nsf.gov/pubs/2015/nsf15505/nsf15505.htm, [accessed 2015-09-01].

[2] Y. Ohmura and Y. Kuniyoshi, "Humanoid robot which can lift a 30kg box by whole body contact and tactile feedback," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2007, pp. 1136–1141.

[3] R. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensing: From humans to humanoids," *Robotics, IEEE Transactions on*, vol. 26, no. 1, pp. 1–20, Feb 2010.

[4] A. Parmiggiani, M. Maggiali, L. Natale, F. Nori, A. Schmitz, N. Tsagarakis, J. S. Victor, F. Becchi, G. Sandini, and G. Metta, "The design of the icub humanoid robot," *Int. Jour. of Humanoid Robots*, 2011.

[5] Y. Ohmura, Y. Kuniyoshi, and A. Nagakubo, "Conformable and scalable tactile sensor skin for curved surfaces," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, May 2006, pp. 1348–1353.

TABLE IV: Quality and Time of Signal Reconstruction with Various Number of Iterations using $\frac{N}{3}$ Measurements

Scenario (No. of Elements)	No. of Measurements	No. of Iterations	Mean Time	Reconstructed Signal Mean PSNR	Reconstructed Signal PSNR Range	Noisy Signal Mean PSNR	Noisy Signal PSNR Range
LIFT BOX (1600)	533	10	10.8	52.5	37.0 - 60.8	46.5	39.1 - 48.7
		20	20.7	52.5	40.9 - 60.8		
		30	30.9	52.5	40.9 - 60.8		
LIFT WRENCH (1600)	533	10	10.6	52.7	26.7 - 60.4	47.2	38.4 - 48.8
		20	20.6	52.7	30.7 - 60.4		
		30	30.1	52.7	31.1 - 60.4		
DRAG WRENCH (1600)	533	10	10.7	47.7	26.3 - 59.8	46.4	37.4 - 48.7
		20	20.7	47.7	28.4 - 59.8		
		30	28.1	47.7	28.5 - 59.8		
PATH (1600)	533	10	10.3	55.9	49.2 - 60.7	48.0	44.8 - 49.1
		20	18.1	55.9	49.1 - 60.7		
		30	28.8	55.9	49.2 - 60.7		
LIFT BOX (4096)	1365	10	8.1	49.9	32.1 - 54.1	46.8	36.6 - 48.5
		20	21.6	49.9	38.9 - 54.1		
		30	25.1	49.9	39.0 - 54.1		
DRAG WRENCH (4096)	1365	10	8.0	46.4	23.1 - 53.9	46.4	36.2 - 48.3
		20	17.0	46.4	29.3 - 53.9		
		30	31.7	46.4	29.8 - 53.9		

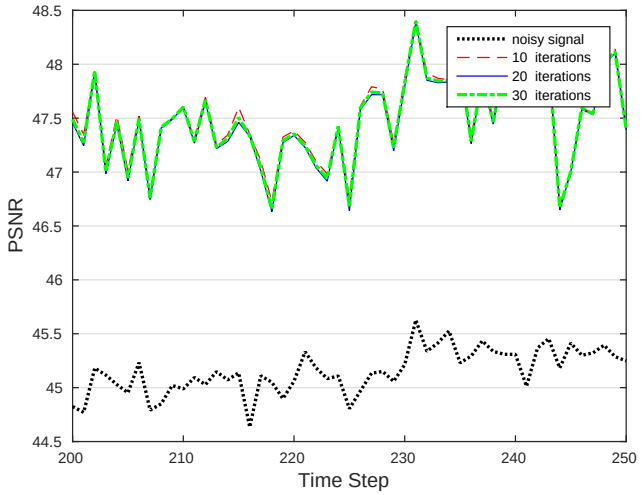

(a) Representative Case Time Sequence - LIFT BOX 64×64

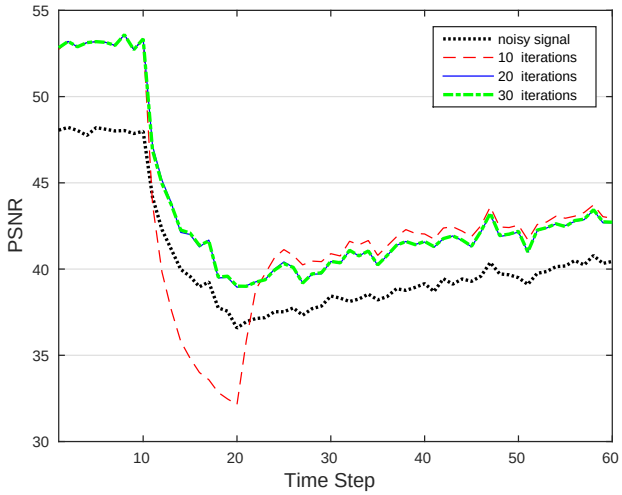

(b) Worst Case Time Sequence - LIFT BOX 64×64

Fig. 4: Both plots are representative time sequences of PSNR for LIFT BOX 64×64 scenario with 10, 20, and 30 iterations and the PSNR of the noisy signal. (a) is a representative time sequence for the majority of time-steps across the scenarios, while (b) is the worst-case time-steps.

[6] A. Schmitz, P. Maiolino, M. Maggiali, L. Natale, G. Cannata, and G. Metta, "Methods and technologies for the implementation of large-scale robot tactile sensors," *Robotics, IEEE Transactions on*, vol. 27, no. 3, pp. 389–400, June 2011.

[7] P. Mittendorfer and G. Cheng, "Humanoid multimodal tactile-sensing modules," *Robotics, IEEE Transactions on*, vol. 27, no. 3, pp. 401–410, June 2011.

[8] T. Mukai, M. Onishi, T. Odashina, S. Hirano, and Z. Luo, "Development of the tactile sensor system of a human-interactive robot "RI-MAN"," *IEEE Trans. Robot.*, vol. 24, no. 2, pp. 505–512, Apr 2008.

[9] M. Davenport, M. Duarte, Y. Eldar, and G. Kutynoik, *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.

[10] "Gazebo," http://gazebosim.org, [accessed 2015-09-13].

[11] A. Habib, I. Ranatunga, K. Shook, and D. O. Popa, "Skinsim: A simulation environment for multimodal robot skin," in *Proc. IEEE Conf. Automation Science and Engineering (CASE)*, 2014, pp. 1226–1231.

[12] J. Tropp and S. J. Wright, "Computational methods for sparse solution of linear inverse problems," *Proc. IEEE*, vol. 98, no. 6, pp. 948–958, 2010.

[13] R. Dahiya, P. Mittendorfer, M. Valle, G. Cheng, and V. Lumelsky, "Directions toward effective utilization of tactile skin: A review," *Sensors Journal, IEEE*, vol. 13, no. 11, pp. 4121–4138, Nov 2013.

[14] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Communications on Pure and Applied Mathematics*, vol. 41, no. 7, pp. 909–996, 1988.

[15] E. J. Candès *et al.*, "Compressive sampling," in *Proc. International congress of mathematicians*, vol. 3, 2006, pp. 1433–1452.

[16] L. Gan, T. Do, and T. Tran, "Fast compressive imaging using scrambled block hadamard ensemble," in *Signal Processing Conference, 2008 16th European*, Aug 2008, pp. 1–5.

[17] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM journal on imaging sciences*, vol. 2, no. 1, pp. 183–202, 2009.